

Exploring the Arista 7010T - An Overview

Arista Networks range of 1Gb, 10Gb, 40Gb and 100GbE Ethernet switches are the foundation for many of the worlds largest data centers.

The introduction of the Arista 7010 Series extends the 7000 Series with a new high performance, feature rich, low power and compact 1Gb Ethernet system that incorporates support for many advanced EOS features, scalable forwarding tables, next generation protocols and a reliable system design.

This white paper provides an overview of the system, forwarding architecture, performance, efficiency, reliability and the network design options.

Switch Overview

The 7010T delivers the flexibility to be deployed as both a 1GbE edge switch in leaf and spine designs and as a dedicated management network switch. With broad support for QoS, security, automation and monitoring features the 7010T provides an ideal solution to the challenges of implementing consistent network policy across both 1G and 10G environments, when combined with the Arista fixed configuration 10G switches.

The Arista 7010 Series are purpose built 1/10GbE data center switches in a compact and energy efficient form factor with wire speed layer 2 and layer 3 forwarding, combined with advanced features for software defined cloud networking. The Arista 7010T switches are fixed configuration 1RU systems, supporting 48 ports of 10/100Mb/1Gb RJ45 ports and 4 SFP+ for both 1Gb and 10GbE connections.

Table 1: 7010 Series Overview	
Characteristic	7010T-48
Switch Height (RU)	1 RU
1/10GbE SFP+ Ports	4
RJ45 10/100/1000Mb Ports	48
Maximum System Throughput (Gbps)	176Gbps
Maximum Forwarding Rate (PPS)	132Mpps
Latency (RJ45 to uplinks)	3usec
Packet Buffer Memory	4MB
Typical Power	52W

The 7010 Series at a Glance



Figure 1: Arista 7010 Series front and back

While many modern networks are migrating rapidly to 10GbE for new servers and implementing 40GbE for advanced IP storage systems, 1GbE Ethernet servers continue to be used widely in both virtualized and traditional non-virtualized environments. This requires flexible networking solutions that deliver features for QoS, security, automation, management and monitoring consistently for both the 1GbE and 10GbE environments. The Arista 7010 Series offers a purpose built high performance and power efficient solution for high density data center deployments. Built on the exact same EOS binary as the entire Arista portfolio the 7010 Series are a natural extension to the 7050X Series bringing the same open platform for automation and DevOps to 1GbE environments.

The 7010 Series comprises two models, the 7010T-48 and the 7010T-48-DC providing a choice of AC or DC power. The 7010T are both well suited to be deployed at the server edge of 1Gb Ethernet leaf and spine designs, where scale out designs using both Layer2 and Layer3 ensures simple deployment. Equally the 7010T can be implemented for integrated lights-out (iLO) or out-of-band (OOB) networks as a dedicated management switch.

The Arista 7010T series delivers line rate, local switching at layer 2 and layer 3, that enables simpler network designs that lowers both the network capital and operational expenses. When used in conjunction with the Arista 7000 series of fixed and modular switches it allows networks to scale using either layer-2 active-active or layer-3 equal cost multi-path (ECMP) designs in a high performance and low-latency two-tier network that provides predictable and consistent application performance for both virtualized and non-virtualized servers. The Arista 7010T switches offer low latency of under 3usec and a shared packet buffer that is allocated dynamically to ports that are congested.

Both 7010T models have 48 ports of 10/100/1000Mb RJ45 and 4 SFP+ interfaces for 1G or 10G uplinks and a broad range of transceivers and cables. Consuming less than 52W the 7010T are extremely power efficient at 0.3W per Gigabit of performance. The Arista 7010 Series are designed for high availability from both a software and hardware perspective with redundant power and cooling, together with EOS high availability features. Built in power redundancy and customer reversible redundant fans allows

cooling airflow in either forward or reverse direction in a single system.

The Arista 7010T runs the same Arista EOS software as all Arista products, simplifying network administration and release certification. Arista EOS is a modular switch operating system with a unique state sharing architecture that cleanly separates switch state from protocol processing and application logic. Built on top of a standard Linux kernel, all EOS processes run in their own protected memory space and exchange state through an in-memory database. This multi-process state sharing architecture provides the foundation for in-service-software updates and self-healing resiliency. Advanced monitoring and automation capabilities such as Zero Touch Provisioning, VM Tracer and Linux based tools can be run natively on the switch with the powerful x86 CPU subsystem.

The 7010 Architecture

At a high level the 7010 series are considered “single chip” switches. The term single chip switch refers to a System on a Chip (SoC) solution, where all hardware forwarding actions are controlled by a single chip. Advances in switching ASIC technology enables the 7010 Series to deliver the number of line rate interfaces delivered in a single chip, while still maintaining high throughput and low power usage.

All stages of the packet forwarding pipeline are performed entirely in the hardware/data plane. The forwarding pipeline is a closed system integrated on the packet processor (PP) of the SoC. The packet processor is capable of providing both the ingress and egress forwarding pipeline stages for all packets that arrive.

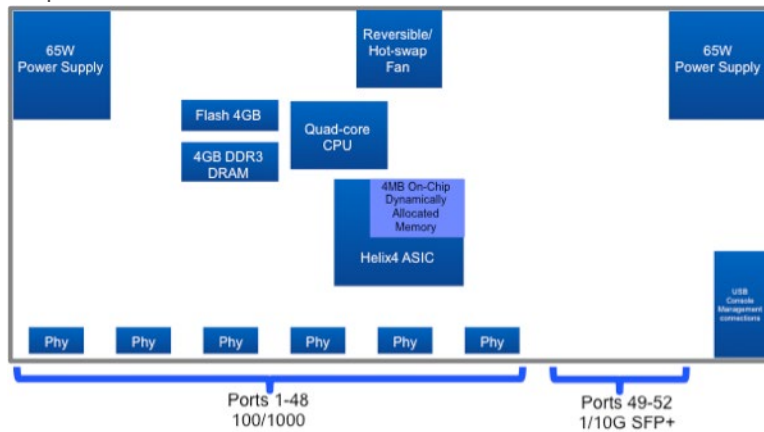


Figure 2: Arista 7010T Architecture

The 7010T delivers 176Gbps of throughput which is sufficient for all 48 ports of 1G and the 4 ports of 10G simultaneously. The maximum packet forwarding rate is also line rate, at up to 132 Million packets per second, for all types of traffic with local processing of L2 and L3 switching and routing, which provides the optimum architecture for top of rack switching. All packet processing for address learning, L2 and L3 lookup, QoS assignment, ingress and egress ACLs and packet rewrites are performed in hardware without any requirement to offload to separate daughter cards or external devices.

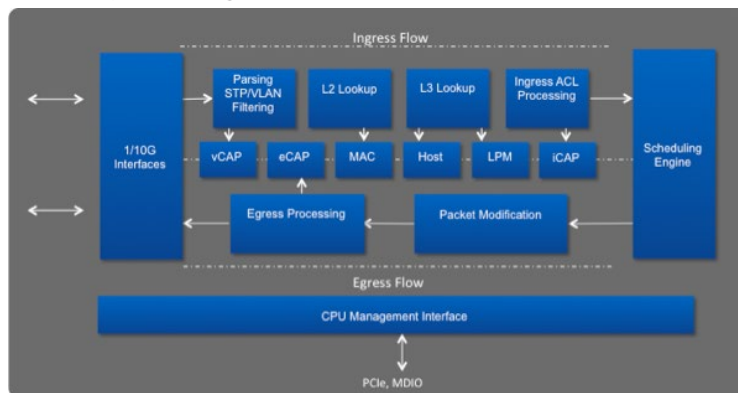


Figure 3: Arista 7010 Series Packet Processing Pipeline

Scaling the Data Plane

In addition to reducing the power and increasing the features available on fixed configuration platforms, the 7010 series also makes significant improvements to both forwarding table density and flexibility. While traditional switches statically allocate resources to specific functions such as MAC Address tables, or IPv4 Host routes, recognizing that no two deployments are identical the 7010 supports a more flexible approach.

Cloud network scalability is directly impacted by the size of a switches forwarding tables. In many systems discrete fixed sized tables are used for each of the common types of forwarding entries, with additional memory required to expand any one table. In cloud environments correctly sizing these tables is critical to allow for flexible network scaling. The Arista 7010 leverages a common unified forwarding table (UFT) for the L2 MAC, L3 Host and IP Multicast forwarding table entries that allows each table to be sized correctly for the solution. Each of these table sizes varies depending on the network deployment scenario. A large choice of pre-defined configurations profiles are available on the 7010 to enable optimal resource utilization for common network topologies and accommodate network and server virtualization. The chart below shows the different table size options with the available UFT modes.

Table 2: Arista 7010 Series Table Scale with UFT

UFT Mode (2 is default)	1	2	3	4	5	
MAC Address Table	84K	64K	44K	24K	4K	Shared Resources
IPv4 Host Routes	4K	24K	44K	64K	84K	
IPv4 Multicast (S,G)	2K	12K	22K	32K	42K	
IPv6 Host Routes	2K	12K	22K	32K	42K	
IPv4 Routes – Unicast	16K					
IPv6 Routes – Unicast	8K (/64) 4K (/128)					
ECMP	64-Way, 1K Groups					
ACLs	4K Ingress / 1K Egress					
IGMP Groups	Up to 8K with 1K unique groups					

Scaling the Control Plane

The CPU on the Arista 7010 Series is used exclusively for control-plane and management functions; all data-plane forwarding logic occurs at the packet processor/Port ASIC level as discussed earlier.

Arista EOS®, the control-plane software for all Arista switches executes on multi-core x86 CPUs with 4 gigabytes of DRAM. As EOS is multi-threaded, runs on a Linux kernel and is extensible, the large RAM and fast multi-core CPUs provide for operating an efficient control plane with headroom for running 3rd party software, either within the same Linux instance as EOS or within a guest virtual machine.

Out-of-band management is available via both a serial console port and the 10/100/1000 Ethernet management interface. The 7010 Series also offer a USB2.0 interface that can be used for a variety of functions including the transferring of images or logs.

Performance

Under all traffic conditions of layer2, or layer 3 IPv4 or IPv6 and multicast testing the 7010T consistently demonstrates 100% line rate performance up to 132Mpps. The charts below show the results of testing.

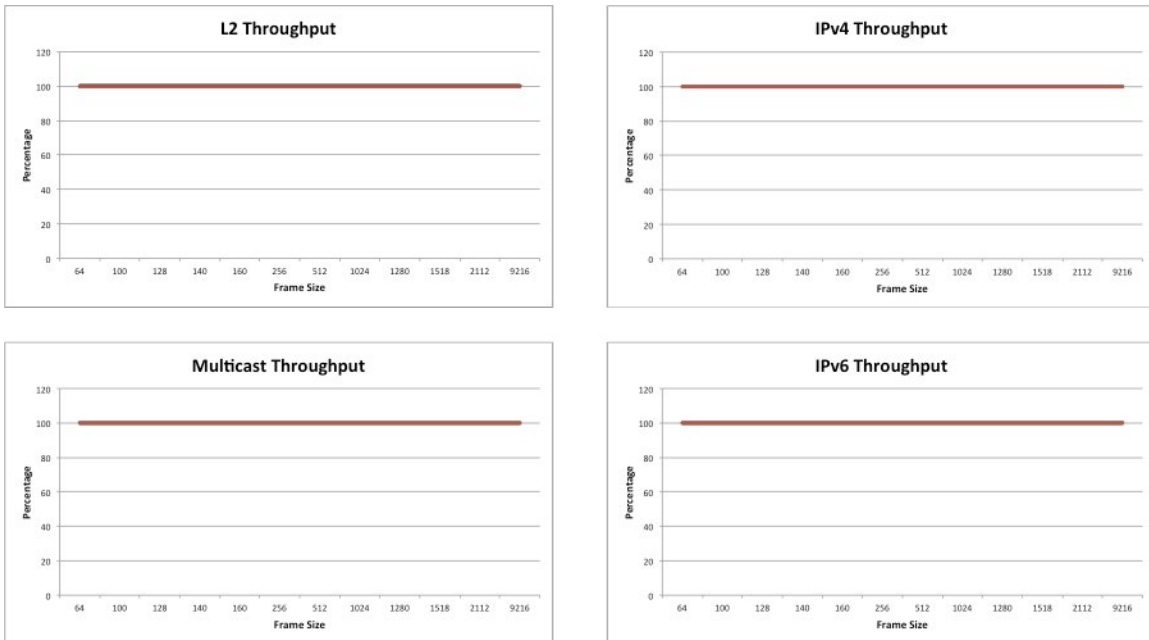


Figure 4: Arista 7010T Performance

The 7010T is designed to meet the needs of both demanding network traffic patterns and to address the requirements for modern SDN and programmability.

Power Efficient and Environmentally Friendly

One aspect of the Arista 7010T that helps it stand alone is the approach to power efficiency and environmental impact.

From the ground up the 7010T was designed for power efficiency. The combination of low power draw and efficient data center focused cooling design ensures a low power footprint. Front to rear or rear to front airflow configurations further avoid the need for specialized cooling infrastructure commonly required for 1G products with side to rear or side to side airflow.

The 7010T requires just 52W for typical use which is less than 1W per port. Comparing to similar products a typical gigabit switch will require over 100W and as much as 200W, more than doubling the power per port. A typical 1RU server will require at least 300W and as much as 700W for maximum power meaning that the 7010T is about 1/10th the power of the servers it networks.

Airflow on the 7010 Series operates as either front to rear, or rear to front, with no side exhaust. Air intakes are located at the rear and front, and a single removable fan tray is used to direct air through the system keeping critical components operating at optimum temperatures. Removing and reversing the fan tray direction changes the airflow direction, avoiding any need to replace fans or an entire system.




Figure 5: Arista 7010T Reversible Airflow

In common with other Arista fixed configuration switches the 7010 implements dynamic cooling with variable speed adjustable fans based on a proportional-integral-derivative (PID) control feedback algorithm that dynamically adjusts the fan speed based on temperature sensors on the inlet, exhaust and all critical components. This ensures that the system is cooled to the correct setting, and does not waste power unnecessarily.

To minimize the environmental impact the 7010T is designed to be as small as possible and reduce packaging. Continuous efforts to reduce the size of packaging, means more can be fit into the shipping containers, which means fewer boats and planes are used resulting in fewer CO2 emissions. It's a minor change but it has a positive impact on the environment. Less packaging also results in less materials for recycling or going to land fill.

For larger quantity customer shipments the 7010T is shipped in high density packaging which is up to 75% more efficient. Four individual switches are contained in a multi-switch package reducing the amount of materials and saving space, which again reduces the environmental impact of shipping and recycling.

Table 3: Arista 7010 Series Power Efficient and Environmentally Friendly

Characteristic	7010T-48
Weight	9.5 lbs (4.3kg)
7010 System Dimensions	19 x 1.75 x 10" (48.3x 4.4 x 25.4cm)
Packaging Dimensions (single system)	21.5 x 16 x 6.5" (54.6 x 40.6 x 16.5cm)
Packaging Dimensions (multiple system)	19.9 x 15.75 x 17.2" (50.5 x 40 x 43.6cm)
Typical Power	52W
BTU / hr	177.4 BTU / hr
CFM (cubic feet per minute)	10.9 (ΔT of 15)
Litre / minute	308 (ΔT of 15)

High Availability

The Arista 7010 Series are designed for high availability from both a software and hardware perspective.

Redundant Power Supplies – The 7010T Series feature two internal redundant, load-sharing AC or DC power supplies to maintain uninterrupted operations. Thanks to the compact design the 7010T Series requires significantly less power than solutions with less performance and features.

Hot Swappable Fan Tray – featuring two fans providing sufficient cooling even if one of the fans were to fail. The reversible fan module provides both front to rear and rear to front airflow in a single unit. The fan module can be removed for replacement or to reverse the direction non-disruptively.

Reversible Fan Module – can be rotated to cool the system in either forward or reverse directions. This simplifies deployment and allows a single system to be configured for either direction. The fan direction can be changed without powering off or restarting the system.

Color coded fan tray – to indicate airflow direction the fan tray uses color coded labels that indicate either front-to-rear or rear-to-front airflow direction.

Live Software Patching – allows the application of critical software patches without requiring a system restart. This also reduces the need to upgrade software for unplanned reasons.

Self Healing Software with Stateful Fault Repair – in the event of a system software fault processes are automatically restarted without affecting other processes and losing any system state.

Multi-chassis LAG (MLAG) – for active / active L2 multi-pathing that eliminates blocked links due to spanning tree.

ECMP Routing – for load balancing and redundancy over layer 3 using open routing protocols – OSPF, BGP, ISIS at both IPv4 and IPv6.

High Reliability – Low component count and low power consumption result in a higher mean time between failure (MTBF).

Scale Out Network Designs at 1G

Based on the Arista universal cloud network design for 2-tier networks, scalable solutions can be delivered using a choice of layer 2, layer3, or with VXLAN to complement L2-over-L3 solutions. Each of the design choices has their advantages and disadvantages. These are discussed in detail in the Arista Scale-Out Cloud Network Designs White Paper.

An important consideration for deploying both layer 2 and layer 3 designs are the forwarding table sizes, and ensuring that the tables offer sufficient capacity to match the design goals. The Arista 7010T uses a unified forwarding table (UFT) that ensures a flexible L2 or L3 design can be implemented without the forwarding table introducing limitations on scale. The diagram below shows how the selection of a L2 or L3 design affects table requirements at the leaf.




Figure 6: Layer 2 and Layer 3 Designs Contrasted

Two-tier leaf-spine topologies can be built at layer 2 or layer 3. If the designs are layer 2, MLAG provides an L2 network that runs active/active with no blocked links, which requires an MLAG peer-link between the spine switches.

It may also be desirable to use MLAG on the leaf switches to connect servers/storage in an active/active manner. In this case, a pair of leaf switches would be an MLAG pair and would have an MLAG peer-link between them. The MLAG peer-link can be a relatively small number of physical links (at least 2) as MLAG prioritizes network traffic so that it remains local to a switch for dual-attached devices.




Figure 7: Spine/Leaf Layer 2 network design for 1G attached nodes with uplinks at 10G

Designs can scale out using layer 3 with up to as many as 32 spine switches in an ECMP layout allowing for a very large fan out of leaf switches. A network may evolve from 2 spine switches to 4, 8, 16 and eventually as many as 32 spine switches. All paths between spine and leaf run active/active utilizing standard routing protocols like BGP and OSPF and up to 32-way ECMP are used to run all paths in active/active mode. The Arista 7010T has a maximum of 4 10G SFP+ interfaces and the following diagrams demonstrate how 1G server designs scale with 4-way ECMP (each leaf switch has 4x10G uplinks for 48x1G server/storage connectivity):




Figure 8: Spine/Leaf layer 3 network design using 4-way ECMP to Arista 7010T Leaf switches

LOM/IPMI Networks

Dedicated OOB networks connect the server IPMI or LOM ports to a single set of networking devices providing management access. As IPMI or LOM ports are always on interfaces if they are unprotected it can allow for complete access to servers and potential malicious activity. The challenge for implementing Out Of Band (OOB) networks is to deliver a simple dedicated network that provides security, reliability and access control services in a cost effective, non-complex and operationally consistent solution to the main IP network.

A set of guidelines for OOB networks includes:

- Use dedicated out of band network devices
- Apply policies to prevent unauthorized access to servers
- Restrict network traffic using access lists and filtering
- Enable monitoring using sFlow or tools such as Splunk to detect malicious behavior
- Secure the OOB switch itself using AAA and security features such as using restricted IP addresses for management or a management vrf.

The 7010T provides a full set of EOS features, including comprehensive security, monitoring and management isolation that allows a secure OOB network to be established. A single EOS binary enables the same operational best practices to be leveraged on all networks. The 7010T power efficiency, small footprint and wire speed performance make it ideal for connecting always on server management interfaces.

Connection Choice and Flexibility

The 7010T offers a broad set of connection capabilities, for use as a high performance 1G system in both legacy 100Mb or virtualized 1G solutions.

All 48 of the RJ45 ports feature auto-negotiation for both 10Mb, 100Mb and 1G for servers, desktop systems or legacy devices. In addition the SFP+ ports allow for dual speed connection at both 1G using a choice of SFP's or at 10G using a wide range of copper and fiber SFP+ that support up to 10km over single-mode fiber. Optics available are fully compliant to IEEE specifications which allows for easy interoperability.

To provision additional network capacity to servers that are virtualized its common to take advantage of link aggregation, or multi-chassis link aggregation (MLAG), to bond multiple 1G interfaces, using LACP or 802.3ad. This works well once the server is configured but does rely on both ends being capable of supporting the LACP features. If a server needs to load its set of software images from a central server using PXE boot or a similar capability it may not support LACP in the default state. The Arista 7010T in line with the Arista 7000 Series supports a LACP fallback feature designed to enable servers to connect in a non-LACP mode to PXE boot, and once the full set of images are loaded and the network interfaces are configured to enable the LACP modes and use link aggregation to the network.

As the Arista 7010T doesn't rely on non-standard forwarding architectures or stacking schemes it can be deployed transparently with multiple vendors at both the spine or leaf tier and will interoperate using open spanning tree or routing protocols. This flexibility avoids any need for fork-lift upgrades to introduce and migrate from existing solutions.

Arista EOS Foundation Features

Arista's scale-out cloud network designs are underpinned on a number of foundation features of Arista's award-winning Extensible Operating System:

Multi Chassis Link Aggregation (MLAG)

MLAG enables devices to be attached to a pair of Arista switches (an MLAG pair) with all links running active/active. MLAG eliminates bottlenecks, provides resiliency and enables layer 2 links to operate active/active without wasting 50% of the bandwidth as is the case with STP blocked links. L3 Anycast Gateway (Virtual ARP / VARP) with MLAG enables an L3 gateway to operate in active/active mode without the overhead of protocols like HSRP or VRRP.

To a neighboring device, MLAG behaves the same as standard link aggregation (LAG) and can run either with Link Aggregation Control Protocol (LACP) or in a static 'mode on' configuration.

The MLAG pair of switches synchronize forwarding state between them such that the failure of one node doesn't result in any disruption or outage as there are no protocols to go from standby to active, or new state to learn as the devices are operating in active/active mode.

MLAG ISSU provides operators with a method to alternate system upgrades without needing to schedule a service outage. MLAG ISSU takes advantage of the rapid rehashing of aggregated links, generally measured in milliseconds, and to synchronize the system database (SysDB) and allowing divergent EOS versions to be run on the two MLAG peers. In short, ISSU allows code upgrades migrating from one version to a new version and adding features and capabilities without causing outages or packet loss.

Zero Touch Provisioning (ZTP)

ZTP enables switches to be physically deployed without pre-staging and for the configuration to be applied automatically without operator intervention. This simplifies deployment, eliminates the risk of misconfiguration and allows network engineering resources to be used for more productive tasks.

New configuration files, updates or changes can be deployed from the central ZTP server using scripts, automated jobs or manual push. Configuration files are automatically backed up, and in the event of a replacement reapplied.

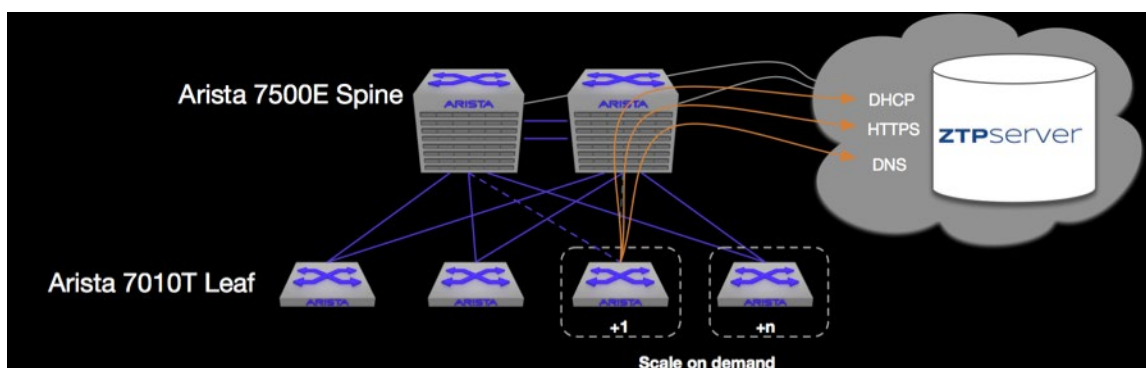


Figure 8: Spine/Leaf layer 3 network design using 4-way ECMP to Arista 7010T Leaf switches

An extension to ZTP, Zero Touch Replacement (ZTR) enables switches to be replaced, with the replacement switch picking up the same image and configuration as the switch it replaced.

VM Tracer

As virtualized data centers have grown in size, the physical and virtual networks that support them have also grown in size and complexity. Virtual machines connect through virtual switches and then to the physical infrastructure, adding a layer of abstraction and complexity. Server side tools have emerged to help VMware administrators manage virtual machines and networks, however equivalent tools to help the network administrator resolve conflicts between physical and virtual networks are less widely available.

Arista VM Tracer provides this bridge by automatically discovering which physical servers are virtualized (by talking to VMware vCenter APIs), what VLANs they are meant to be in (based on policies in vCenter) and then automatically apply physical switch port configurations in real time with vMotion events. This results in automated port configuration and VLAN database membership and the dynamic adding/removing VLANs from trunk ports.

VM Tracer also provides the network engineer with detailed visibility into the VM and physical server on a physical switch port while enabling flexibility and automation between server and network teams.

Automated Configuration, Visibility and Management

In line with the Arista EOS automation and provisioning tools the 7010T leverages a range of features that simplify deployment, orchestration, configuration management, automation and performance monitoring.



Figure 10: Integrated Automation with DevOps toolkits and EOS eAPI

- ZTP Server allows the use of configuration templates and scripts to automate deployment and eliminate human error associated with initial configuration of systems and ongoing configuration changes.
- Networks automation integration with tools such as Chef, Puppet and Ansible ensure that existing DevOps tools are leveraged to minimize repetitive operations.
- Support for AEM EventMonitor, Splunk integration, and VMTracer provide visibility into network behavior.
- Integration with Microsoft OMI, VMWare, Openstack and common SNMP based management systems for provisioning and network compliance management.
- CloudVision is an Arista provided EOS software extension that enables an XMPP client to interact with multiple Arista switches in the topology providing a single management plane.

Arista EOS: A Platform for Scale, Stability and Flexibility

Arista Extensible Operating System, or EOS®, is the most advanced network operating system in the world. It combines modern-day software and O/S architectures, transparently restartable processes, open platform development, a Linux kernel, and a stateful publish/subscribe database model.

At the core of EOS is the System Data Base, or SysDB for short. SysDB is machine generated software code based on the object models necessary for state storage for every process in EOS. All inter-process communication in EOS is implemented as writes to SysDB objects. These writes propagate to subscribed agents, triggering events in those agents. As an example, when a user-level ASIC driver detects link failure on a port it writes this to SysDB, then the LED driver receives an update from SysDB and it reads the state of the port and adjusts the LED status accordingly. This centralized database approach to passing state throughout the system and the automated way SysDB code is generated reduces risk and error, improving software feature velocity and provides flexibility for customers who can use the same APIs to receive notifications from SysDB or customize and extend switch features.




Figure 11: Arista EOS Software Architecture showing some of the Agents

Arista's software engineering methodology also benefits customers in terms of quality and consistency:

- Complete fault isolation in the user space and through SysDB effectively convert catastrophic events to non-events. The system self-heals from more common scenarios such as memory leaks. Every process is separate, no IPC or shared memory fate sharing, endian-independent, and multi-threaded where applicable.
- No manual software testing. All automated tests run 24x7 and with the operating system running in emulators and on hardware Arista scales protocol and unit testing cost effectively.
- Keep a single system binary across all platforms. This improves the testing depth on each platform, improves time-to-market, and keeps feature and bug compatibility across all platforms.

EOS, and at its core SysDB, provide a development framework that enables the core concept - Extensibility. An open foundation, and best-in-class software development models deliver feature velocity, improved uptime, easier maintenance, and a choice in tools and options.

Conclusion

The 7010 Series combined with the Arista leaf-spine design methodology and the Arista 7000 Family provides flexible design solutions for network architects looking to deliver market-leading performance driven networks, which scale from several hundred hosts all the way up several hundred thousand hosts.

The 7010 Series delivers 48 x 1GbE and 4 1/10GbE SFP+ ports designed specifically to operate at line-rate in real world deployments. With up to 176Gbps of forwarding capacity, and 84K MAC and IPv4 hosts with UFT the 7010 Series provides the port density, table scale, feature set and forwarding capacity essential in today's data center environments.

The 7010T are power efficient, deliver wire speed layer 2 and layer 3 using hardware forwarding, are designed for high availability and for simple integration to existing designs using either MLAG or ECMP. With simplified configuration management and integration with automation frameworks the challenges of supporting discrete devices is eliminated whilst preserving the best of the flexibility of simple, topology independent, open networking.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. July 01, 2018 02-0010-02